

EV 316937433

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

**Security Policy Update Supporting At Least One
Security Service Provider**

Inventor(s):

Thomas F. Fakes

Anders M. E. Samuelsson

ATTORNEY'S DOCKET NO. MS1-1704US

1 **TECHNICAL FIELD**

2 This invention relates to security, and more particularly to security policy
3 update supporting at least one security service provider.
4

5 **BACKGROUND**

6 Computers have become increasingly commonplace and increasingly
7 interconnected, such as by the Internet and the use of electronic mail (email).
8 Unfortunately, such increased interconnectivity has resulted in increased attacks
9 against computers by malicious users. During such attacks, malicious users
10 oftentimes try to introduce malicious programs on to other users' computers (e.g.,
11 designed to disable the other users' computers, obtain information from the other
12 users' computers, launch attacks against still other computers, and so forth),
13 attempt to disable a computer so that its performance is greatly impaired (e.g., by
14 bombarding the computer with requests), and so forth. From the perspective of
15 the user of the attacked computer, these attacks can range from being annoying to
16 devastating, potentially resulting in the computer being inoperable or confidential
17 information being copied from the computer.

18 In order to combat these attacks, security services that run on a computer
19 have been created to protect that computer from attack. Examples of such security
20 services include antivirus programs and firewall programs. One aspect of these
21 security services is that they typically need to be updated regularly in order to be
22 able to protect against the latest attacks.

23 However, one problem that can exist when updating security services is that
24 the capabilities of the services may be hindered during the updating process. As
25 such, the computer can face an increased vulnerability to attack when being

1 updated. Thus, it would be beneficial to have a way to reduce the vulnerability of
2 computers during the updating of the security services on the computer.

3 4 **SUMMARY**

5 Security policy update supporting at least one security service provider is
6 described herein.

7 In accordance with certain aspects, each of one or more security service
8 providers receives a set of new rules to be enforced as part of a new security
9 policy. Each security service provider processes the new rules in order to be ready
10 to begin using the new rules, but continues to use the previous set of rules until
11 instructed to begin using the new rules. When all of the one or more security
12 service providers are ready to begin using the new rules, they are instructed to
13 begin using the new rules at which point all of the security service providers begin
14 using the set of new rules substantially concurrently.

15 16 **BRIEF DESCRIPTION OF THE DRAWINGS**

17 The same numbers are used throughout the document to reference like
18 components and/or features.

19 Fig. 1 illustrates an example environment in which the security policy
20 update can be used.

21 Fig. 2 is a flowchart illustrating an example process for updating a security
22 policy.

23 Fig. 3 is a flowchart illustrating another example process for updating a
24 security policy.

25 Fig. 4 illustrates an example of a general computer environment.

DETAILED DESCRIPTION

Updating a security policy in one or more security service providers is described herein. Multiple security service providers can be employed in a device to provide security, such as antivirus protection, intrusion detection, intrusion prevention, and so forth. A security policy can be used to describe the rules that are to be applied by these providers. When changes are made to security policy, updated rules are supplied to the various security service providers, and these various security service providers change over to start using the updated rules at substantially the same time.

Fig. 1 illustrates an example environment 100 in which the security policy update can be used. Environment 100 includes a source device 102 and a device 104. Source device 102 stores or otherwise has access to or can obtain one or more security policies 106. A security policy 106 describes how various security service providers on device 104 should operate. Although only one source device 102 is illustrated in Fig. 1, it is to be appreciated that a device 104 can obtain security policies (or different portions of security policies) from multiple different source devices 102. Similarly, although only one device 104 is illustrated in Fig. 1, it is to be appreciated that multiple devices 104 can obtain security policies from the same source device 102.

Device 104 can be any of a variety of devices where a security policy may be enforced, and source device 102 can be any of a variety of devices that make a security policy (or portion of a security policy) available. Examples of devices 104 and source devices 102 include desktop or workstation computing devices, server computing devices, portable or handheld computing devices, game

1 consoles, network appliances, cellular phones, personal digital assistants (PDAs),
2 networking devices (e.g., routers, gateways, firewalls, wireless access points, etc.),
3 and so forth.

4 Environment 100 can represent any of a variety of a settings, such as
5 networks in home, business, educational, research, etc. settings. For example,
6 source device 102 may be a server device on a corporate LAN, and device 104
7 may be a desktop or portable computing device on the corporate LAN. By way of
8 another example, source device 102 may be a server device on the Internet, and
9 device 104 may be a desktop computing device at a user's home.

10 Device 104 includes a policy reader module 110, a manager module 112, a
11 rule set generator module 114, dynamic rules data store 116, and one or more
12 security service providers (also referred to as security engines) 118. It is to be
13 appreciated that although separate modules are shown in device 104 of Fig. 1,
14 alternatively one or more of these modules may be combined into a single module,
15 and/or one or more of these modules may be separated into two or modules.

16 Generally, to update the security policy being enforced by security engines
17 118 on device 104, policy reader 110 obtains a security policy 106 from source
18 device 102. Rule set generator 114 uses the newly obtained security policy 106 to
19 generate, for each of the various security engines 118, a set of one or more rules
20 and associated data. These sets of rules are then communicated to the various
21 security engines 118, and the associated data is stored in dynamic rules data store
22 116. The associated data can also be communicated to the security engines 118.
23 Upon receiving the set of one or more rules, each security engine 118 processes
24 the new set of rules, getting ready to begin using the new set of rules. However,
25 each security engine 118 continues to use its current set of rules until instructed by

1 manager 112 to change to the new set. Manager 112 instructs all of the security
2 engines 118 to change to the new set of rules after manager 112 receives an
3 indication from each of the security engines 118 that it is ready to change to the
4 new set of rules.

5 Each security engine performs various actions to help secure device 104
6 from malicious users and/or malicious programs. Such malicious users and/or
7 malicious programs may attempt to disable device 104 or disable functionality
8 (e.g., programs) of device 104, obtain data from device 104 (e.g., passwords or
9 other confidential information), use device 104 (e.g., to assist in attacking other
10 devices), and so forth. A security engine 118 can be any service that assists in
11 protecting against such malicious users and/or malicious programs.

12 Examples of security engines 118 include, but are not limited to: an
13 antivirus engine, a firewall engine, an intrusion detection engine, a vulnerability
14 analysis engine, and a behavioral blocking engine. An antivirus engine includes
15 functionality to protect device 104 from being infected by viruses, worms, Trojan
16 horses, and so forth. A firewall engine includes functionality to protect device 104
17 from being accessed over a network connection by other undesired devices (e.g.,
18 any device that a user or administrator of device 104 does not want accessing
19 device 104), and also functionality to block malicious code from propagating to
20 other devices over the network. An intrusion detection engine includes
21 functionality to identify when a malicious program and/or user has accessed
22 device 104 (e.g., due to the antivirus and/or firewall engines having been
23 overridden or having failed, etc.) and take appropriate action (e.g., notify a user or
24 administrator, attempt to disable the malicious program or halt the malicious user's
25 access, etc.). A vulnerability analysis engine includes functionality to attempt to

1 detect vulnerabilities in device 104 (e.g., due to security engines that have not
2 been installed or updated correctly, security engines or other components being
3 configured incorrectly, and so forth). A behavior blocking engine includes
4 functionality to monitor programs running on device 104 and detect improper
5 behavior (e.g., improper network or storage device access) by those programs.

6 In certain embodiments, manager module 112 coordinates the updating of
7 security policy in device 104. Manager 112 receives the indications from the
8 various security engines 118 that indicate the security engines 118 are ready to
9 change to the new set of rules, and gives an indication to the security engines 118
10 when they should begin using the new set of rules.

11 Policy reader module 110 obtains a new security policy from source device
12 102. Policy reader module 110 may be configured to check whether a new
13 security policy is available from source 102 at regular or irregular intervals, or
14 alternatively may receive an indication from some other component (e.g., manager
15 112, source device 102, some other device not shown in Fig. 1, and so forth) that it
16 should obtain a new security policy from source 102 (or check whether a new
17 security policy is available from source 102). Policy reader 110 may identify to
18 source 102 a particular security policy that reader 110 desires to obtain, or
19 alternatively may merely request the most recent security policy for device 104
20 from source 102. A comparison between the current security policy being used by
21 device 104 and the most recent security policy may be made to determine whether
22 the most recent security policy is already being enforced on device 104. Such a
23 comparison could be made by source 102, policy reader 110, or alternatively some
24 other component.
25

1 When the new security policy is obtained from source 102, rule set
2 generator 114 generates a set of rules for each of the different security engines
3 118. Different security engines may use different rules when enforcing the
4 security policy on device 104. For example, one security engine 118 may be a
5 firewall whereas another security engine 118 may be an antivirus component. The
6 security policy may identify rules that are specific to the antivirus engine (and thus
7 the firewall engine need not be concerned with them), and may also identify rules
8 that are specific to the firewall engine (and thus the antivirus engine need not be
9 concerned with them).

10 In certain embodiments, the security policy itself is a list of rules and
11 associated data. The security policy may also include an indication of which rules
12 and data are for which security engines, or alternatively no such indication may be
13 included (e.g., relying on device 104 to determine which rules are for which
14 security engines). The security policy allows designers to have a single record or
15 file of all the rules involved in the protection of device 104, without having to
16 separate the rules across different records or files for the different security engines.

17 Additionally, using the techniques described herein, new security policies
18 can be prepared by designers that shift responsibility for protecting against
19 particular attacks from one security engine to another. For example, protection
20 against a particular type of attack may be enforced by an antivirus program in one
21 security policy but changed to being enforced by a firewall program in a new
22 security policy. Using the techniques described herein, the designers can be
23 confident that this shift in responsibility will occur in all of the security engines
24 substantially concurrently, thereby reducing the vulnerability of the device 104 to
25 attacks during the shift.

1 Rule set generator 114 identifies, based on the security policy, which rules
2 and associated data (if any) are used by which of the security engines 118. For
3 each security engine 118, rule set generator 114 generates a set of rules for that
4 security engine and makes that generated set of rules available to that security
5 engine (e.g., the set of rules may be transmitted or sent to the security engine, the
6 security engine may be informed of a location in memory where the generated set
7 of rules can be obtained, etc.). This generation can be performed in a variety of
8 different manners. For example, a new set of rules may be generated by generator
9 114 without regard for the current rules being enforced by the security engines
10 118. By way of another example, the current set of rules may be modified or
11 changed to incorporate any differences between the current and new set of rules.
12 Additionally, rule set generator 114 may simply copy the rules from the security
13 policy, or alternatively rule set generator 114 may generate the rules based on
14 information in the security policy that describes the rules.

15 In certain embodiments, the security policy identifies which rules are to be
16 distributed to which security engines. For example, each rule may be associated
17 with a particular label or identifier (e.g., Security Engine (1), or Antivirus engine,
18 etc.). Rule set generator 114 can thus use these identifiers in generating the sets of
19 rules for the various security engines 118. In alternate embodiments, rule set
20 generator 114 may infer which rules are to be distributed to which security
21 engines. In other embodiments, a combination of these techniques may be used
22 (e.g., for some rules the security policy may identify which security engine they
23 are to be assigned to, and for other rules the security policy generator 114 may
24 infer which security engine they are to be assigned to).

1 The set of rules generated by rule set generator 114 can take any of a
2 variety of different forms. In certain embodiments, the rules follow an if-then
3 structure. Using this structure, the rule defines a particular condition(s) and a
4 corresponding particular action(s) or result(s). During enforcement of the rule, if
5 that particular condition(s) is detected then the corresponding particular action(s)
6 or result(s) is performed. Any of a variety of conditions and corresponding results
7 can be identified by a rule. Examples of particular conditions include: attempts to
8 access particular resources (e.g., memory locations, network addresses or ports,
9 other programs, files on a storage device, and so forth), attempts to write data to
10 particular locations (e.g., to particular memory locations, to particular locations on
11 a storage device, etc.), attempts to run particular programs, various aspects of the
12 current operating state of the device 104 (e.g., resources available, programs
13 running, etc.), and so forth. Examples of particular results include: preventing a
14 resource from being accessed, preventing data from being written to particular
15 locations, preventing a program from being run, generating a notification that the
16 occurrence of the condition in the rule was detected (e.g., recording its occurrence
17 in a log, sending a message to a user or other computer, and so forth). The
18 particular results can also be permissive in nature rather than preventive. For
19 example, the results could indicate that a particular resource or location can be
20 accessed only if the condition in the rule is satisfied by the device 104, or that a
21 particular program can be run only if the condition in the rule is satisfied by the
22 device 104.

23 In certain embodiments, device 104 includes dynamic rules data store 116
24 which is the data associated with the various rules being enforced by security
25 engines 118. In certain embodiments, dynamic rules data store 116 may include

1 two sets of data: one set for the current rules being enforced by security engines
2 118, and another set for the new rules that security engines 118 are being updated
3 to enforce. When a new security policy is received, generator 114 updates
4 dynamic rules data store 116 with the data associated with the sets of new rules
5 passed to security engines 118.

6 Each security engine 118 includes a rule change module 120 that receives a
7 set of one or more rules from rule set generator 114. The data associated with the
8 rules may be received from generator 114 along with the rules, or alternatively
9 rule change module 120 may obtain the data it desires from dynamic rules data
10 116. Additionally, it should be noted that although rule set generator 114 is
11 discussed above as generating a set of rules for each security engine 118 based on
12 the security policy, alternatively each security engine 118 may receive the entire
13 security policy (or most of the security policy) and generate their own set of rules
14 rather than receiving the set from generator 114.

15 Rule change module 120 processes the new set of rules as needed in order
16 to generate new internal rules 122 that enforce the new policy. The processing of
17 the new set of rules to generate new internal rules 122 refers to whatever actions
18 are necessary for the security engine to take in order to place the new set of rules
19 in a state that they can be enforced by the security device. For example, this
20 processing may include converting the new set of rules to an internal format,
21 storing rules in particular memory locations, organizing rules into a particular
22 arrangement or order, etc. Rule change module 120 may generate new rules in
23 any of a variety of manners; module 120 may keep the rules in the same format as
24 they were received from generator 114 or alternatively convert the rules to an
25 internal format use by the security engine.

1 Regardless of how the new rules are generated, each security engine 118
2 maintains a current set of rules 124 which enforce the previous security policy for
3 device 104 (the security policy which is being updated). While generating the new
4 rules 122, and even after the new rules 122 are generated, the security engine
5 continues to enforce current rules 124. Security engine 118 does not begin
6 enforcing new rules 122 until instructed to do so (e.g., by manager 112).

7 After rule change module 120 has finished generating new rules 122,
8 module 120 indicates to manager 112 that it has finished and is ready to switch to
9 using the new rules (and thus begin enforcing the new security policy). After
10 manager 112 has received such an indication from all of the security engines 118,
11 manager 112 instructs each of the security engines 118 to begin using the new
12 rules. Manager 112 waits to instruct each of the security engines to begin using
13 the new rules until after manager 112 receives the indication from all of the
14 security engines 118. Once instructed to do so, each security engine 118 begins
15 using the new rules 122. As soon as a security engine 118 begins using the new
16 rules 122, it can delete the rules it was previously using (current rules 124).

17 In some situations, a security engine 118 may fail in processing the new
18 rules. In such situations, the security engine 118 returns an indication of such
19 failure to manager 112. Alternatively, manager 112 may impose a time limit on
20 responses from the security engines 118. If all security engines 118 have not
21 responded with an indication that they are ready to begin using the new rules
22 within the time limit, then manager 112 can assume that one or more of the
23 security engines 118 has failed in processing the new rules.

24 When manager 112 identifies that one or more of the security engines 118
25 has failed in processing the new rules, manager 112 does not instruct any of the

1 security engines 118 to begin using the new rules. Rather, manager 112 sends an
2 indication to abort the changeover to the new rules (this may also be referred to as
3 a rollback). Such an abort or rollback indication informs each of the security
4 engines 118 that it is to ignore the new rules received from generator 114 as well
5 as any new rules 122 resulting from its processing, and continue to use the current
6 rules. In certain embodiments, the security engines 118 can safely delete the new
7 rules they generated (or are in the process of generating) in response to such an
8 abort or rollback indication.

9 In certain embodiments, each security engine 118 waits until it can nearly
10 ensure that it can begin using the new rules before informing manager 112 that it
11 is ready to begin using the new rules. In other words, the security engine 118
12 waits to inform manager 112 that it is ready to begin using the new rules until the
13 security engine 118 is to the point in processing the new rules that it is virtually
14 impossible for the security engine 118 to fail to begin enforcing those rules when
15 instructed to do so. In certain embodiments, this is accomplished by security
16 engine 118 generating the new set of rules 122, and maintaining a pointer to which
17 of the rule sets (set 122 or set 124) it is to use. After the new set of rules 122 is
18 generated, security engine 118 indicates to manager 112 that the security engine is
19 ready to begin using the new set of rules. Then, when instructed to begin using the
20 new set of rules, the security engine can simply change its pointer from the old set
21 of rules to the new set of rules. The security engine 118 can nearly ensure that it
22 can change its pointer and begin using the new rules. It is to be appreciated that
23 "nearly ensure" does not require a 100% guarantee that failure is absolutely
24 impossible. It is possible that certain situations could still arise that would result
25

1 in failure (e.g., a power loss or virus attack that prohibits changing of the pointer).

2 However, it is also to be appreciated that the chances of failure are very small.

3 Manager 112 can instruct the security engines 118 to begin using the new
4 set of rules (also referred to as switching over to the new set of rules) in any of a
5 variety of different manners. The manner that is used, however, should operate to
6 inform all of the security engines at substantially the same time so that all of the
7 security engines can begin using their new sets of rules at substantially the same
8 time (also referred to herein as substantially concurrently). By having all of the
9 security engines begin using their new sets of rules at substantially the same time,
10 any vulnerability of device 104 due to the rule changeover is reduced. Generally,
11 the closer in time that the security engines begin using their new sets of rules, the
12 lesser the vulnerability of the device during the changeover to the new set of rules.
13 Following are some examples of ways in which manager 112 can instruct the
14 security engines 118 at substantially the same time to begin using their new sets of
15 rules.

16 One way in which manager 112 can instruct the security engines 118 to
17 begin using the new set of rules is to use an event object that can be fired across all
18 of the security engines at once. For example, each security engine, upon receipt of
19 the new rules from rule set generator 114, sets an internal flag to start polling a
20 particular event each time the rules are accessed (during its normal operation of
21 protecting device 104). Manager 112 can then instruct the security engines to
22 begin using their new sets of rules by firing the event (the same one being polled
23 by the security engines). So, after the event is fired, any subsequent polling of the
24 event will reflect that the event has been fired and thereby inform the polling
25 security engine that the new rule set should be used. For example, in response to

1 detecting the event having been fired, the pointer in the security engine can be
2 changed to point to the new set of rules.

3 In addition to polling the event, a thread may also be run by the security
4 engine that waits on the event. When the event is fired, the thread detects the
5 firing so that the security engine is informed that the new rule set should be used.
6 For example, in response to the thread detecting that the event has fired, the
7 pointer in the security engine can be changed to point to the new set of rules.

8 Once the event has been fired and the new set of rules is being used, the
9 security engine can stop polling the event. Additionally, if a thread waiting on the
10 event is run by the security engine, that thread can be terminated.

11 Another way in which manager 112 can instruct the security engines 118 to
12 begin using the new set of rules is to call a function exposed by each of the
13 security engines 118 (e.g., a "switch" function). Calling such a function of a
14 security engine 118 instructs that security engine to begin using the new set of
15 rules. For example, in response to such a function being invoked on a security
16 engine 118, the security engine changes its pointer to point to the new set of rules.

17 Another way in which manager 112 can instruct the security engines 118 to
18 begin using the new set of rules is to notify each of the security engines 118 of a
19 shared data structure that each security engine 118 can access. Manager 112 can
20 inform each security engine of the shared data structure at different times, such as
21 by calling a function on each security engine 118 (e.g., an "identify data structure"
22 function), or by identifying the shared data structure when the new rules are
23 passed to the security engine 118. The shared data structure can take any of a
24 variety of different forms, such as a location in memory (e.g., in random access
25

1 memory (RAM) or a nonvolatile memory such as Flash memory), a file stored on
2 a local or remote storage device, and so forth.

3 Each security engine checks this shared data structure (e.g., each time the
4 rules are accessed (during its normal operation of protecting device 104)) to
5 determine its value. Manager 112 can instruct each of the security engines 118 to
6 begin using the new rule set by changing the value(s) stored in the shared data
7 structure. For example, the shared data structure may initially store a value of
8 "previous" or a value of 0 to indicate that the current set of rules are to be used,
9 and when it is time to switch to begin using the new rule set manager 112 can
10 write a value of "new" or "switch" or a value of 1 to the shared data structure to
11 indicate that the new set of rules are to be used.

12 As discussed above, dynamic rules data store 116 stores the data associated
13 with the various rules being enforced by security engines 118. As such, when
14 device 104 is being updated to begin enforcing a new policy, the data used by the
15 security engine 118 may also change. This data can also change during the
16 operation of device 104 (e.g., a security engine 118 may later request data from or
17 store data in dynamic rules data 116 store). In order for the proper data to be made
18 available to the security engines 118, when updating the security policy dynamic
19 rules data store 116 may operate in the same manner as a security engine 118.
20 That is, two sets of rules data would be maintained – the first set would be used
21 prior to the switch and the second set would be used after the switch. The new
22 data would be stored in data store 116, and data store 116 would return an
23 indication to manager 112 when it is ready to begin using the new set of data.
24 Data store 116 then continues to use the previous set of data until receiving an
25 instruction from manager 112 to begin using the new set of data.

1 It should be noted that the various components in device 104 illustrated in
2 Fig. 1 can be implemented within the same application process executing on
3 device 104. For example, policy reader 110, manager 112, dynamic rules data
4 116, rule set generator 114, and security engines 118 may all be part of the same
5 application process.

6 Alternatively, different components in device 104 illustrated in Fig. 1 can
7 be implemented across two or more application processes executing on device
8 104. For example, one or more security engines 118 may run in a process that is
9 separate from the other security engines 118 as well as separate from policy reader
10 110, manager 112, dynamic rules data 116, and rule set generator 114. Allowing
11 different components to run in different application processes allows, for example,
12 different developers to design different plug-in components (e.g., different plug-in
13 security engines) to enhance the security of device 104. These additional plug-in
14 components would be upgraded to enforce new policies in the same manner as
15 other non-plug-in components.

16 When separating the components across multiple application processes, a
17 mechanism is used to allow the various components to communicate with one
18 another. This communication allows, for example, sets of new rules and data to be
19 passed to security engines in different processes, data to be passed from security
20 engines in different processes to dynamic rules data 116, instructions to begin
21 using the new sets of rules to be passed to security engines in different processes,
22 and so forth. By way of example, the components of Fig. 1 may be implemented
23 as Component Object Model (COM) components. Additional information
24 regarding the Component Object Model architecture is available from Microsoft
25 Corporation of Redmond, Washington.

1 It should be noted that in the discussions herein, each security engine is
2 instructed to begin using its new set of rules by manager 112. Alternatively, this
3 instruction may be implemented in other manners that still allow each security
4 engine to begin using its new set of rules substantially concurrently. For example,
5 rather than using manager 112, a control mechanism to instruct each security
6 engine to begin using its new set of rules may be distributed across the various
7 security engines 118. This could be accomplished, for example, by each of the
8 security engines notifying each other security engine that it is ready to begin using
9 the new set of rules, with none of the security engines beginning to use its new set
10 of rules until all of the security engines have notified all of the other security
11 engines that they are ready to begin using the new set of rules.

12 Fig. 2 is a flowchart illustrating an example process 200 for updating a
13 security policy. Process 200 is implemented by a component(s) that is
14 coordinating the updating of the security policy on a device, such as manager 112
15 of device 104 of Fig. 1. Process 200 may be performed in software, hardware,
16 firmware, or combinations thereof.

17 Initially, a new policy to be enforced on the device is obtained (act 202).
18 The policy may be obtained in a "pull" manner, where device 104 initiates the
19 access to the source of the new policy to check whether a new policy is available
20 from the source. The policy may alternatively be obtained in a "push" manner,
21 where device 104 is informed of (e.g., sent a message or other indication of) the
22 availability of a new security policy or of the new security policy itself.

23 Regardless of how the new policy is obtained, once the new policy is
24 obtained a new set of rules and/or data associated with the rules for the new policy
25 is provided to each of the security devices (act 204). As discussed above, different

1 sets of rules and/or data can be generated based on the new policy for each
2 security engine.

3 Return values are then received from the security engines (act 206). In
4 certain implementations, each security engine returns, to the component
5 implementing process 200, a value signifying "OK" or a value signifying "Fail".
6 When a security engine returns a value signifying "OK" it indicates that the
7 security engine has processed the set of rules and/or data that it received and is
8 ready to begin using the new set of rules and/or data. For example, all that
9 remains is for the security engine to change its pointer to point to the new set of
10 rules rather than the previous set of rules. However, when a security engine
11 returns a value signifying "Fail", it indicates that the security engine could not (or
12 did not) process the set of rules and/or data and that the security engine is not able
13 to begin using the new set of rules and/or data. Additionally, as discussed above a
14 time limit (also referred to as a timeout value or a threshold amount of time) may
15 be imposed on responses from security engines – if a security engine does not
16 respond with a value signifying "OK" or "Fail" within this time limit the
17 component implementing process 200 treats the security engine as if it had
18 returned a value signifying "Fail".

19 It is to be appreciated that the sending of rules and the receiving of
20 responses (acts 204 and 206) is an asynchronous process. Different security
21 engines may take different amounts of time to process the rules and/or data they
22 receive, and process 200 simply waits until all of the security engines have
23 finished their respective processing (up to any optional time limit that is imposed).

24 Process 200 proceeds based on whether all of the security engines have
25 returned a value signifying "OK" (act 208). If all of the security engines have

1 returned a value signifying "OK", then all of the security engines are ready to
2 begin using the new set of rules, so all of the security engines are instructed to
3 begin using the new set of rules (act 210).

4 However, if at least one of the security engines does not return a value
5 signifying "OK", then a rollback call is issued to each security engine (act 212).
6 This rollback call essentially aborts the update process, so none of the security
7 engines will begin to use the new set of rules yet (even those security engines that
8 had returned a value signifying "OK").

9 Fig. 3 is a flowchart illustrating another example process 300 for updating a
10 security policy. Process 300 is implemented by a security engine on a device,
11 such as a security engine 118 on device 104 of Fig. 1. Process 300 may be
12 performed in software, hardware, firmware, or combinations thereof.

13 Initially, a new set of rules and/or data are received for the new policy to be
14 enforced (act 302). As discussed above, the rules and data may be received at
15 substantially the same time, or alternatively the security engine may obtain data
16 from a data store (e.g., data store 116 of Fig. 1) as needed. The new rules and/or
17 data are then processed (act 304). Processing of the new rules and/or data creates
18 an internal set of rules for the security engine to follow (e.g., in an internal format
19 of the security engine) in enforcing the new security policy.

20 Process 300 proceeds based on whether the processing of the rules was
21 successful (act 306). If the security engine has finished processing the set of rules
22 and/or data that it received and is ready to begin using the new set of rules and/or
23 data (e.g., all that remains is for the security engine to change its pointer to point
24 to the new set of rules rather than the previous set of rules), then the processing
25 was successful. Otherwise, the processing was not successful. If the processing

1 was successful than a value signifying "OK" is returned (act 308). However, if the
2 processing was not successful then a value signifying "Fail" is returned (act 310).
3 The return values in acts 308 and 310 are returned to a component(s) that is
4 coordinating the updating of the security policy on the device (e.g., manager 112
5 of Fig. 1).

6 Regardless of the value returned, the security engine continues to use the
7 previous or old set of rules until instructed to rollback or begin using new rules
8 (act 312). If instructed to begin using the new rules, then the security engine
9 begins using the new rules and/or data (act 314), such as by changing a pointer
10 from its previous set of rules to its new set of rules. The instruction to begin using
11 the new rules can be received by the security engine in any of a variety of
12 manners, as discussed above.

13 However, if instructed to rollback, then the security engine discards any
14 results of processing the new rules and/or data (act 316). This discarding can be
15 performed regardless of whether the security engine has finished processing the
16 new set of rules it received.

17 Thus, as can be seen in Fig. 3, the security engine continues to use its
18 previous set of rules until an indication to switch to the new rules is received. By
19 the time such an indication is received, the security engine is ready to begin using
20 the new rules, and very little time is required for the switch to occur. For example,
21 the security engine may simply need to switch a pointer to point to its new set of
22 rules rather than its previous set of rules.

23 Fig. 4 illustrates an example of a general computer environment 400, which
24 can be used to implement the techniques described herein. The computer
25 environment 400 is only one example of a computing environment and is not

1 intended to suggest any limitation as to the scope of use or functionality of the
2 computer and network architectures. Neither should the computer environment
3 400 be interpreted as having any dependency or requirement relating to any one or
4 combination of components illustrated in the example computer environment 400.

5 Computer environment 400 includes a general-purpose computing device in
6 the form of a computer 402. Computer 402 can be, for example, a source 102 or a
7 device 104 of Fig. 1. The components of computer 402 can include, but are not
8 limited to, one or more processors or processing units 404, a system memory 406,
9 and a system bus 408 that couples various system components including the
10 processor 404 to the system memory 406.

11 The system bus 408 represents one or more of any of several types of bus
12 structures, including a memory bus or memory controller, a peripheral bus, an
13 accelerated graphics port, and a processor or local bus using any of a variety of
14 bus architectures. By way of example, such architectures can include an Industry
15 Standard Architecture (ISA) bus, a Micro Channel Architecture (MCA) bus, an
16 Enhanced ISA (EISA) bus, a Video Electronics Standards Association (VESA)
17 local bus, and a Peripheral Component Interconnects (PCI) bus also known as a
18 Mezzanine bus.

19 Computer 402 typically includes a variety of computer readable media.
20 Such media can be any available media that is accessible by computer 402 and
21 includes both volatile and non-volatile media, removable and non-removable
22 media.

23 The system memory 406 includes computer readable media in the form of
24 volatile memory, such as random access memory (RAM) 410, and/or non-volatile
25 memory, such as read only memory (ROM) 412. A basic input/output system

1 (BIOS) 414, containing the basic routines that help to transfer information
2 between elements within computer 402, such as during start-up, is stored in ROM
3 412. RAM 410 typically contains data and/or program modules that are
4 immediately accessible to and/or presently operated on by the processing unit 404.

5 Computer 402 may also include other removable/non-removable,
6 volatile/non-volatile computer storage media. By way of example, Fig. 4
7 illustrates a hard disk drive 416 for reading from and writing to a non-removable,
8 non-volatile magnetic media (not shown), a magnetic disk drive 418 for reading
9 from and writing to a removable, non-volatile magnetic disk 420 (e.g., a "floppy
10 disk"), and an optical disk drive 422 for reading from and/or writing to a
11 removable, non-volatile optical disk 424 such as a CD-ROM, DVD-ROM, or other
12 optical media. The hard disk drive 416, magnetic disk drive 418, and optical disk
13 drive 422 are each connected to the system bus 408 by one or more data media
14 interfaces 426. Alternatively, the hard disk drive 416, magnetic disk drive 418,
15 and optical disk drive 422 can be connected to the system bus 408 by one or more
16 interfaces (not shown).

17 The disk drives and their associated computer-readable media provide non-
18 volatile storage of computer readable instructions, data structures, program
19 modules, and other data for computer 402. Although the example illustrates a hard
20 disk 416, a removable magnetic disk 420, and a removable optical disk 424, it is to
21 be appreciated that other types of computer readable media which can store data
22 that is accessible by a computer, such as magnetic cassettes or other magnetic
23 storage devices, flash memory cards, CD-ROM, digital versatile disks (DVD) or
24 other optical storage, random access memories (RAM), read only memories
25 (ROM), electrically erasable programmable read-only memory (EEPROM), and

1 the like, can also be utilized to implement the example computing system and
2 environment.

3 Any number of program modules can be stored on the hard disk 416,
4 magnetic disk 420, optical disk 424, ROM 412, and/or RAM 410, including by
5 way of example, an operating system 426, one or more application programs 428,
6 other program modules 430, and program data 432. Each of such operating
7 system 426, one or more application programs 428, other program modules 430,
8 and program data 432 (or some combination thereof) may implement all or part of
9 the resident components that support the distributed file system.

10 A user can enter commands and information into computer 402 via input
11 devices such as a keyboard 434 and a pointing device 436 (e.g., a “mouse”).
12 Other input devices 438 (not shown specifically) may include a microphone,
13 joystick, game pad, satellite dish, serial port, scanner, and/or the like. These and
14 other input devices are connected to the processing unit 404 via input/output
15 interfaces 440 that are coupled to the system bus 408, but may be connected by
16 other interface and bus structures, such as a parallel port, game port, or a universal
17 serial bus (USB).

18 A monitor 442 or other type of display device can also be connected to the
19 system bus 408 via an interface, such as a video adapter 444. In addition to the
20 monitor 442, other output peripheral devices can include components such as
21 speakers (not shown) and a printer 446 which can be connected to computer 402
22 via the input/output interfaces 440.

23 Computer 402 can operate in a networked environment using logical
24 connections to one or more remote computers, such as a remote computing device
25 448. By way of example, the remote computing device 448 can be a personal

1 computer, portable computer, a server, a router, a network computer, a peer device
2 or other common network node, and the like. The remote computing device 448 is
3 illustrated as a portable computer that can include many or all of the elements and
4 features described herein relative to computer 402.

5 Logical connections between computer 402 and the remote computer 448
6 are depicted as a local area network (LAN) 450 and a general wide area network
7 (WAN) 452. Such networking environments are commonplace in offices,
8 enterprise-wide computer networks, intranets, and the Internet.

9 When implemented in a LAN networking environment, the computer 402 is
10 connected to a local network 450 via a network interface or adapter 454. When
11 implemented in a WAN networking environment, the computer 402 typically
12 includes a modem 456 or other means for establishing communications over the
13 wide network 452. The modem 456, which can be internal or external to computer
14 402, can be connected to the system bus 408 via the input/output interfaces 440 or
15 other appropriate mechanisms. It is to be appreciated that the illustrated network
16 connections are examples and that other means of establishing communication
17 link(s) between the computers 402 and 448 can be employed.

18 In a networked environment, such as that illustrated with computing
19 environment 400, program modules depicted relative to the computer 402, or
20 portions thereof, may be stored in a remote memory storage device. By way of
21 example, remote application programs 458 reside on a memory device of remote
22 computer 448. For purposes of illustration, application programs and other
23 executable program components such as the operating system are illustrated herein
24 as discrete blocks, although it is recognized that such programs and components
25

1 reside at various times in different storage components of the computing device
2 402, and are executed by the data processor(s) of the computer.

3 Various modules and techniques may be described herein in the general
4 context of computer-executable instructions, such as program modules, executed
5 by one or more computers or other devices. Generally, program modules include
6 routines, programs, objects, components, data structures, etc. that perform
7 particular tasks or implement particular abstract data types. Typically, the
8 functionality of the program modules may be combined or distributed as desired in
9 various embodiments.

10 An implementation of these modules and techniques may be stored on or
11 transmitted across some form of computer readable media. Computer readable
12 media can be any available media that can be accessed by a computer. By way of
13 example, and not limitation, computer readable media may comprise "computer
14 storage media" and "communications media."

15 "Computer storage media" includes volatile and non-volatile, removable
16 and non-removable media implemented in any method or technology for storage
17 of information such as computer readable instructions, data structures, program
18 modules, or other data. Computer storage media includes, but is not limited to,
19 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
20 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
21 tape, magnetic disk storage or other magnetic storage devices, or any other
22 medium which can be used to store the desired information and which can be
23 accessed by a computer.

24 "Communication media" typically embodies computer readable
25 instructions, data structures, program modules, or other data in a modulated data

1 signal, such as carrier wave or other transport mechanism. Communication media
2 also includes any information delivery media. The term “modulated data signal”
3 means a signal that has one or more of its characteristics set or changed in such a
4 manner as to encode information in the signal. By way of example, and not
5 limitation, communication media includes wired media such as a wired network or
6 direct-wired connection, and wireless media such as acoustic, RF, infrared, and
7 other wireless media. Combinations of any of the above are also included within
8 the scope of computer readable media.

9 One or more flowcharts are described herein and illustrated in the
10 accompanying Figures. The ordering of acts in these flowchart(s) are examples
11 only – these orderings can be changed so that the acts are performed in different
12 orders and/or concurrently.

13 Although the description above uses language that is specific to structural
14 features and/or methodological acts, it is to be understood that the invention
15 defined in the appended claims is not limited to the specific features or acts
16 described. Rather, the specific features and acts are disclosed as exemplary forms
17 of implementing the invention.